METHODS AND APPARATUS FOR USING MULTIPLE REASSEMBLY MEMORIES FOR PERFORMING MULTIPLE FUNCTIONS

Field of the Invention

The present invention relates generally to packet processing systems, and more particularly to the use of multiple reassembly memories for performing multiple functions associated with such packet processing systems.

Background of the Invention

As is known, during certain processes performed in a router or other type of packet switch of a packet processing system, packets may be segmented into subsets or portions of data referred to as "cells." For example, packets may be segmented into cells during router framing operations. However, these cells of data must be reassembled back into packets or protocol data units (PDUs) for use by other processes or functions performed in the router.

Conventional routers typically reassemble packets and store them in a common reassembly memory for subsequent use by multiple functions performed by the router. Such functions may include, for example, packet classification and packet scheduling. However, as is known, reassembling packets for use by such multiple functions requires very high input and output bandwidth.

The use of a common, high bandwidth memory to perform multiple functions has many significant drawbacks. First, such a high bandwidth memory can be quite expensive. It also typically causes any associated memory interface device to be expensive too because of requirements such as extra pins, special buffers and special control mechanisms. Further, the use of a common, high bandwidth reassembly memory makes the memory partitioning design task very difficult, particularly if the design is implemented in multiple integrated circuits.

It is therefore apparent that a need exists for techniques which address these and other drawbacks associated with the use of a common, high bandwidth memory for storing reassembled packets for subsequent use in multiple functions performed in a packet processing system.

Summary of the Invention

The present invention provides packet processing techniques which employ multiple reassembly memories for performing multiple functions associated with a packet processing system thereby avoiding the drawbacks attributable to the conventional use of a common, high bandwidth memory.

In one aspect of the invention, a processing system comprises first processing circuitry for performing a first function, and first memory circuitry coupled to the first processing circuitry for storing received packets, wherein at least a portion of the packets stored by the first memory circuitry are usable by the first processing circuitry in accordance with the first function. The processing system further comprises at least second processing circuitry for performing a second function, and at least second memory circuitry coupled to the second processing circuitry for storing at least a portion of the same packets stored in the first memory circuitry, wherein at least a portion of the packets stored in the second memory circuitry are usable by the second processing circuitry in accordance with the second function.

Thus, the invention provides a packet processing system such that the memory required to perform the first and at least second functions is separately partitioned into a first memory and at least a second memory which respectively provide enough bandwidth to store the same data, or at least a subset of required data, and to allow performance of the corresponding function.

It is to be understood that the first processing circuitry, the first memory circuitry, the second processing circuitry and the second memory circuitry may be implemented on the same integrated circuit. Alternatively, the first processing circuitry and the first memory circuitry may be implemented on a first integrated circuit, and the second processing circuitry and the second memory circuitry are implemented on a second integrated circuit.

The processing system may also comprises first reassembly circuitry, coupled to the first memory circuitry, for reassembling subsets of received packets prior to storing the packets in the first memory circuitry, and at least second reassembly circuitry, coupled to the second memory circuitry, for reassembling at least a portion of the same subsets of packets reassembled by the first

reassembly circuitry, prior to storing the packets in the second memory circuitry. The packet subsets are preferably cells.

The processing system may also comprise parsing circuitry, coupled to the first reassembly circuitry and the second reassembly circuitry, for parsing information from the received packets for use by the first reassembly circuitry and the second reassembly circuitry in respectively reassembling the packets.

In an illustrative embodiment, the first processing circuitry and the first memory circuitry comprise a network processor. In such case, the first function may be a packet classifying operation. Further, the second processing circuitry and the second memory circuitry may comprise a traffic manager. In such case, the second function may be a packet scheduling operation.

In another illustrative embodiment, the first processing circuitry and the second processing circuitry operate in a packet switching device such as a router. In such case, the first processing circuitry and the second processing circuitry operate between a packet network interface and a switch fabric of the packet switching device.

Advantageously, the packet processing techniques of the invention increase system performance and reduce system cost due to a reduction in required memory bandwidth associated with performing multiple functions and the associated impact of simplification in the overall processing system design.

These and other objects, features and advantages of the present invention will become apparent from the following detailed description of illustrative embodiments thereof, which is to be read in connection with the accompanying drawings.

Brief Description of the Drawings

FIG. 1 is a block diagram illustrating a packet processing system employing multiple reassembly memories according to an embodiment of the present invention;

FIG. 2 is a flow diagram illustrating a packet processing methodology employing multiple reassembly memories according to an embodiment of the present invention; and

FIG. 3 is a block diagram illustrating a packet processing system employing two reassembly memories for respectively performing packet classification and packet scheduling functions according to an embodiment of the present invention.

Detailed Description of Preferred Embodiments

The present invention will be illustrated below in conjunction with an exemplary packet processing system which includes packet reassembly operations. It should be understood, however, that the invention is more generally applicable to any packet processing system in which it is desirable to avoid the drawbacks attributable to the use of a common, high bandwidth memory.

It is to be understood that the term "processor" as used herein may be implemented, by way of example and without limitation, utilizing a microprocessor, central processing unit (CPU), digital signal processor (DSP), application-specific integrated circuit (ASIC), or other type of data processing device or processing circuitry, as well as portions and combinations of these and other devices or circuitry.

The present invention in an illustrative embodiment avoids the drawbacks attributable to the use of a common, high bandwidth reassembly memory by employing multiple reassembly memories for respectively performing multiple functions associated with multiple processors of a packet processing system. Among other advantages, the invention yields improved system performance and reduced design expense. Also, the invention yields reduced system expense. This is because it is often cheaper to have two memories of a given bandwidth rather than one memory of twice the bandwidth, the latter being what the conventional approach would require.

FIG. 1 shows a packet processing system 100 employing multiple reassembly memories according to an embodiment of the present invention. The system 100 includes a cell parser 102 and N packet processors 104-1 through 104-N, where N may be an integer equivalent to the number of processors that the processing system 100 is designed to support. Each packet processor104 includes a packet or PDU reassembler 106 (106-1 through 106-N) and a PDU memory 108 (108-1 through 108-N). Each processor also has at least one packet-related function 110 (110-1 through 110-N) associated therewith.

It is to be understood that each processor may perform more than one function and that any two processors may perform distinct portions of the same function. Also, each of the N processors may be implemented on N, more than N, or less than N, integrated circuits or processing devices (including one integrated circuit or processing device).

Further, the packet processing system 100 may be implemented in a router or other type of packet switch. In such case, the functions to be performed by the respective processors in accordance with their respective reassembly memories may be, by way of example and without limitation, packet classification, packet scheduling, etc.

Advantageously, as shown, the packet processing system 100 is designed such that the memory required to perform the N functions is partitioned into N memories (108-1 through 108-N) which respectively provide enough bandwidth to reassemble the same data, or at least a selection of required data, to perform the corresponding function. Thus, the data may be reassembled and stored in parallel in each of the N memories. The parallel operations may be simultaneous or substantially simultaneous (e.g., delayed by some amount of time). Thus, the above-described drawbacks associated with a common, high bandwidth reassembly memory are advantageously avoided.

FIG. 2 shows a packet processing methodology 200 employing multiple reassembly memories according to an embodiment of the present invention. Specifically, FIG. 2 will be described below with reference to the multiple reassembly memory arrangement shown in FIG. 1.

As mentioned above, during certain router processes (e.g., router framing operations), packets may be segmented into subsets or portions of data called "cells." Cells are therefore typically subsets of packets and may include a header and a payload. For example, a cell may include the beginning, middle, or end (or some combination thereof) of a full packet. However, all or some of these cells must be reassembled back into packets or PDUs for respective use by functions 110-1 through 110-N.

Thus, in step 202, cells received by the data processing system 100 are parsed by parser 102. The parsing operation includes extracting reassembly information from the cells. As is known, reassembly information is data which instructs the reassemblers 106 how to reassemble the data

associated with a packet into its original order or sequence before the packet was segmented into cells.

In step 204, the reassembly information and the cells are passed from the parser 102 onto each reassembler 106 of each processor 104. The reassemblers then respectively reassemble the cells into PDUs, as is known, using the reassembly information.

It is to be understood that while all reassemblers may reassemble the same data, this is not required. That is, each reassembler may need only reassemble data specific to the function to be performed by its associated processor. Thus, the bandwidth and size of each reassembly memory may be tailored to the processing requirements of the function.

Next, in step 206, the reassembled PDUs are stored in each PDU memory 108 of each processor 104. Lastly, in step 208, the reassembled PDUs stored in each memory 108 (or at least a portion of such stored PDUs) are then utilized by each processor to perform each function 110.

As mentioned above, steps 204 through 208 of FIG. 2 may be performed in each processor simultaneously or substantially simultaneously, depending on the overall design of the packet processing system 100. However, this is not required.

Referring now to FIG. 3, a packet processing system 300 is shown employing two reassembly memories for respectively performing packet classification and packet scheduling functions in accordance with the present invention. More particularly, it is to be understood that FIG. 3 depicts a specific example (where N equals 2) of the processing system 100 of FIG. 1.

The packet processing system 300 includes a network processor 302 with a memory 304 and a traffic manager 306 with a memory 308. It is to be understood that in this exemplary embodiment, the network processor 302 with memory 304 represents one processor or processing circuitry with its corresponding reassembly memory circuitry, while the traffic manager 306 and memory 308 represent the other processor or processing with its corresponding reassembly memory circuitry.

As is known, a network processor such as is shown in FIG. 3 generally controls the flow of packets between a physical transmission medium, such as a physical layer portion of an asynchronous transfer mode (ATM) network or a synchronous optical network (SONET), and a switch fabric in a router or other type of packet switch. One function of a network processor is

packet classification. A traffic manager such as is shown in FIG. 3 generally operates in conjunction with a network processor and performs, among other functions, buffer management and packet scheduling.

Accordingly, as shown, the network processor 302 and the traffic manager 306 are functionally positioned between a network interface 310, which provides an interface (e.g., a physical layer interface and framer) between the processors and a network 312 (e.g., ATM, SONET, etc.) and a switch fabric 314. The network 312 is a network from which packets or other packet data is received. The switch fabric 314 controls switching of packets. The two processors are also responsive to a host CPU 316 which may provide overall control over the two processors.

As is known, routers and switches generally include multiple processors, e.g., arranged in the form of an array of line cards with one or more processors associated with each line card. Thus, it is to be understood that in this embodiment the network processor 302 and the traffic manager 306 may represent processors implemented on a line or port card of a router or other type of packet switch. The network processor and traffic manager may be implemented on the same integrated circuit or different integrated circuits.

Thus, with reference back to the steps of FIG. 2 in view of this particular embodiment, cells in the system 300 (e.g., generated by segmentation associated with the network interface 310 and/or the switch fabric 314) are parsed for reassembly information (step 202). The parser may be implemented in either the network processor 302 or the traffic manager 306. Then, the cells may be simultaneously or substantially simultaneously reassembled (step 204) into original packets by each of the processors 302 and 306 and stored in their associated memories 304 and 308 (step 206). The reassembled packets stored in each memory are then utilized by each processor to perform each function (step 208), e.g., packet classification in the network processor 302 and packet scheduling in the traffic manager 306.

Although illustrative embodiments of the present invention have been described herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various other changes and modifications may be made by one skilled in the art without departing from the scope or spirit of the invention.